

Appl. No. 09/496,990
Amdt. Dated 01/24/2005
Reply to Office action of 10/05/2004

APPENDICES

1. Chapter 3: Traffic Management, CISCO PNNI Reference Guide, March 2001,
http://www.cisco.com/application/pdf/en/us/guest/products/ps525/c2001/ccmigration_09186a0080088c56.pdf.
2. Interface Table - Chapter 4: Configuration Tables and Fields, CISCO WAN Modeling Tools User Guide, November 200,
http://www.cisco.com/application/pdf/en/us/guest/products/ps6106/c2001/ccmigration_09186a0080117a97.pdf.



Traffic Management

The following traffic management functions are supported on the PNNI node:

- Asymmetrical traffic requirements
- Connection admission control (CAC)
- Qbin for AutoRoute and PNNI
- Mapping of SVC/SPVC Traffic Parameters to UPC capabilities
- Dual leaky bucket UPC generic cell-rate algorithm (GCRA 1 and GCRA2)
- ABR EFCI marking, Relative Rate marking, Explicit Rate marking and VS/VD control
- Frame Discard for AAL5

Service Categories

The following service categories as defined in ATM Forum Traffic Management 4.0 are supported:

- Constant bit rate (CBR)
- Real-time VBR (rt-VBR)
- Non-real-time VBR (nrt-VBR)
- Unspecified bit rate (UBR)
- Available bit rate (ABR)

Note ABR is currently only supported in SPVCs. ABR will be supported in SVCs when UNI 4.0 is supported.

Table 3-1 Supported Service Categories

Attribute	CBR	rt-VBR	nrt-VBR	UBR	ABR
PCR, CDVT	Yes	Yes	Yes	Yes	Yes
SCR, MBS	—	yes	yes	—	—
MCR (UNI 4.0)	—	—	—	—	Yes
RAIG CLP (PNNI 1.0)	Yes	Yes	Yes	—	—
peak-to-peak CDV (UNI 4.0)	Yes	Yes	No	No	No

Table 3-1 Supported Service Categories (continued)

Attribute	CBR	rt-VBR	nrt-VBR	UBR	ABR
max CTD (UNI 4.0)	Yes	Yes	No	No	No
CLR (UNI 4.0)	Yes	Yes	Yes	No	Network Specific

Connection Admission Control

Each connection between an ATM CPE (ATM end user) and the WAN switching network establishes a contract with the network based on QoS parameters. This contract specifies the intended traffic flow, including values for peak bandwidth, average sustained bandwidth, and burst size.

The ATM CPE is responsible for adhering to the contract by means of traffic shaping. Traffic shaping is the use of queues to constrain data bursts, limit peak data rate, and smooth jitter so that the traffic complies with the contract.

Nodes have the option of using traffic policing to enforce the networking contract. There are two methods for policing traffic from ATM CPE into the network:

- Connection admission control (CAC)
- Usage parameter control (UPC)

CAC verifies that sufficient network resources are available to accept the call. The desired ATM traffic contract is specified in the Bearer Capability, the Traffic Descriptor, and QoS information elements of the Setup message. These UNI signaling Setup message elements are mapped to UPC parameters in the SES PNNI node.

G-CAC and A-CAC

CAC functions are performed on each individual node. G-CAC (Generic CAC) is performed by the routing protocol when it selects a routing path for a call. The PNNI routing protocol needs to make sure that the routing path it selected will satisfy the basic traffic parameter requirement specified by the call. The checking performed by the routing protocol at border node is called G-CAC. For G-CAC description, refer to "ATM Forum PNNI 1.0 Spec. af-pnni-0055.000". A routing path which satisfied the G-CAC checking may not guarantee the call will have traffic resource on each node of its routing path. This is attributed to the routing data base synchronization delay in the network. To ensure a call has traffic resource it requested on each node of its routing path, a local CAC (A-CAC) must be performed when the call arrived on the node. A call will be released or cranked back if it fails on A-CAC. A-CAC performance is described in the section below.

Both resource-related criteria and policy-related criteria are used for CAC. And the CAC is performed per service category. The CAC is applied to all types of interfaces, UNI, NNI, and TSP.

Resource-based CAC

For resource-related CAC, the CAC algorithm calculates the Equivalent Cell Rate (ECR) of a connection prior to establishing the connection, and ensures high efficiency of network resources in accordance with the following criteria:

- Requested service category
- Requested conformance definition
- Traffic descriptor, PCR, SCR, MCR, and MBS
- CDVT
- Requested UPC conformance
- QoS parameters (namely, CDV, CTD, and CLR)
- Buffer size at each of the QBIN
- Line speed at each port
- Policy-related parameters, minimum and maximum bandwidth, booking factor for a service category

When a call is admitted with the ECR, the requested QoS for the connection is guaranteed even under the worst scenario as described by the traffic descriptors of the connection.

Policy-based CAC

For policy-related CAC, the following parameters are configurable per service category.

- Booking factor per service category per interface
- Guaranteed minimum bandwidth per service category per interface
- Maximum bandwidth allowed per service category per interface
- Guaranteed minimum number of connections per service category per interface
- Maximum number of connections allowed per service category per interface
- Maximum bandwidth allowed for a single connection per service category per interface

Usage Parameter Control

Usage Parameter Control (UPC) may be performed at the ingress of network edge nodes only, or alternatively may be performed at every intermediate node. UPC is enabled on a per interface basis as specified by a service class template on the switches. The traffic contracts are defined in the ATM Forum Traffic Management Specification 4.0 and UNI 3.x, and are summarized in the following table.

Table 3-2 Traffic Contracts

Service Category	CBR	CBR	rt, nrt, VBR	rt, nrt VBR	CBR	rt, nrt VBR	UBR	UBR	ABR
UNI 4.0 Conformance	¹ CBR.2	CBR.3	VBR.2	VBR.3	CBR.1	VBR.1	UBR.1	UBR.2	ABR
Parameters									
PCR (0+1)	x	x	x	x	x	x	x	x	x
PCR (0)	x	x							

Table 3-2 Traffic Contracts (continued)

Service Category	CBR	CBR	rt, nrt, VBR	rt, nrt VBR	CBR	rt, nrt VBR	UBR	UBR	ABR
UNI 4.0 Conform-ance	CBR.2	CBR.3	VBR.2	VBR.3	CBR.1	VBR.1	UBR.1	UBR.2	ABR
CDVT(0+1)	x	x	x	x	x	x	x	x	x
CDVT(0)	x	x							
SCR (0+1)						x			
SCR (0)			x	x					
MBS (0+1)						x			
MBS(0)			x	x					
Tagging		x		x				x	
Best Effort							x	x	
MCR									x
Transmit ATC ²	5	5	9, 10	9, 10	7	19, 11	10	10	12
ILMI Traffic Descriptor Type	NoClpNo oScr	NoClpNo oScr	ClpNoTaggingScr	ClpTaggingScr	NoClpNo oScr	NoClpScr	NoClpNo oScr	NoClpNo oScr	ClpNoTagging-Mer
ILMI Best Effort Indicator	false	false	false	false	false	false	true	true	false
Specification	3.x	3.x	3.x 4.0	3.x 4.0	3.x 4.0	3.x 4.0	3.x 4.0	4.0	4.0

1. This traffic definition applies to the CBR traffic sets specified in TM 3.x.

2. Transmit ATC is defined in TM 4.0, which is not supported in the current release.

Available Bit Rate

Available Bit Rate (ABR) can operate in the following modes:

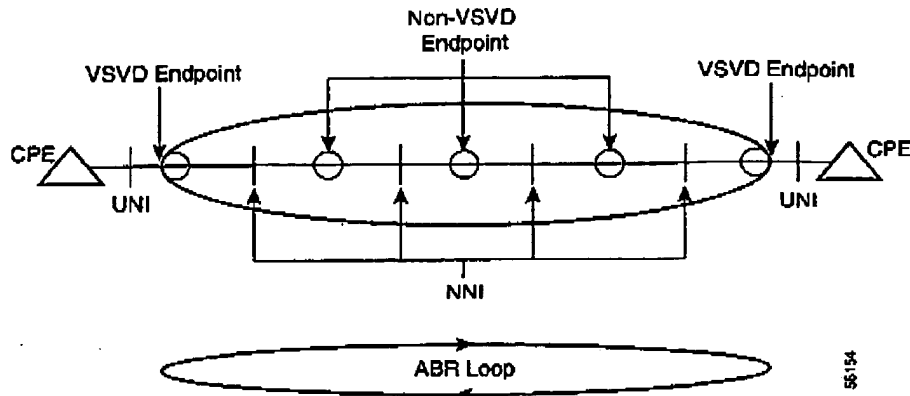
- EFCI marking
- Relative Rate marking
- Explicit Rate marking
- VS/VD control.

EFCI marking and VS/VD control are configured by the service template for an interface. SPVC supports VS/VD on a per connection basis. Relative Rate marking (CI control) and Explicit Rate marking are configured on a per slot basis, and this configuration applies to both AutoRoute and PNNI ABR connections on the slot.

ABR VS/VD Control

The ABR closed loop virtual source/virtual destination (VS/VD) control is set up between the originating and terminating UNIs which are set up as VS/VD endpoints. An interface is provisioned so that ABR endpoints on that interface operate as a VS/VD endpoints.

Figure 3-1 ABR VS/VD Control Loop



The ABR parameters (Table 3-3) are obtained from the following SETUP message information:

- Broadband Bearer Capability Bearer Class must be set to ABR and User Plane Connection Configuration must be point-to-point.
- Traffic Descriptor IE contains the forward and backward MCR and PCR values.
- ABR Setup Parameters Mandatory IE contains the forward and backward ICR, TBE, RIF and RDF values, and the CRM value.
- ABR Additional Parameters Optional IE contains the forward and backward Nrm, Trm, CDP and ADTP parameters.

If ABR parameters are not supplied in the SETUP message then default values are used.

Table 3-3 ABR Signal Parameters

Parameter	Default	Value Range	Units	Description	IE
PCR	mandatory	50 through Line Rate	cells/s	Peak Cell Rate	Traffic Descriptor
MCR	0	50 through Line Rate	cells/s	Minimum Cell Rate	Traffic Descriptor
ICR	PCR	MCR through PCR	cells/s	Initial Cell Rate	ABR Setup Parameters
RIF	1/16	Power of 2: 1/32768 through 1	-	Rate Increase Factor	ABR Setup Parameters
RDF	1/16	Power of 2: 1/32768 through 1	-	Rate Decrease Factor	ABR Setup Parameters
TBE	16777215	0 through 16777215	cells	Transient Buffer Exposure	ABR Setup Parameters
FRTT	-	0 through 16777215	usec	Fixed Round Trip Time	ABR Setup Parameters

Table 3-3 ABR Signal Parameters (continued)

Parameter	Default	Value Range	Units	Description	IE
Nrm	32	Power of 2: 2 through 256	cells	Maximum cells per forward RM-cell	ABR Additional Parameters
ADTF	0.55	0.01-10.23	sec	ACR Decrease Time Factor	ABR Additional Parameters
Trm	100	100 times power of 2: $100 \times 2^7 - 100 \times 2^0$	insec	Time between Forward RM-cells	ABR Additional Parameters
CDF	1/16	Power of 2: 1/64 - 1	-	Cutoff Decrease Factor	ABR Additional Parameters
CRM	TBE/Nrm	1 - 4095	cells	Cumulative RM-cell Count	Computed

When the SES node receives a Setup message, the following processing, which includes Connection Admission Control (CAC), occurs:

- Validate the Bearer Class and Traffic Descriptor combinations.
- Determine the SES PNNI controller Service Category from the Bearer Capability information element and the Traffic Descriptor information element. The SES PNNI controller Service Category is either CBR, rt-VBR, nrt-VBR, ABR, or UBR.
- Verify that the required bandwidth is less than the Available Cell Rate. CAC is performed on the BXM.
- Verify that a channel is available on the port by obtaining a free LCN.
- Verify that a VPCI is available.
- Map the Traffic Parameters to the port UPC capabilities.

Frame Discard for AAL5

PNNI nodes support frame discard for AAL5 on PNNI 1.0 signaling to support SPVC applications. This can be configured per SPVC. Frame discard for SVC is configured on per port basis (on or off). Control of frame discard for individual SVC will be supported in later release when UNI4.0 is supported.

Overbooking



The PNNI Controller supports both overbooking and underbooking when it sets up a port's CAC policy parameters. The CLI to change booking factor is `cnfnpportcac portid service_category -bookfactor utilization-factor`. The Service Category Utilization Factor (SCUF) is in the range of 1 to 200. The default is 100. An SCUF = 1 means the controller book 1% of requested bandwidth (i.e. request 100M BW on a call, the controller records only 1 M of bandwidth is reserved). This is an overbooking. On other hand, an SCUF = 200 means the controller book 200% of requested bandwidth. This is an underbooking.

PNNI uses Generic CAC (GCAC) to select a path with sufficient BW for a newly configured connection.

An ATM interface module performs Actual CAC (ACAC) to check that there are sufficient resources on the interface before adding the connection on the individual interface. If ACAC fails for an interface that passed GCAC, crankback occurs.

Oversubscription results from overbooking of connection BW on a link so that the sum of configured BW for connections exceeds the actual link BW. This is configured using a booking factor of 1 to 99. Oversubscription (a booking value 1 to 99) overbooking.

Undersubscription results from underbooking of connection BW on a link so that the sum of configured BW for connections will always be less than the actual link BW. This is configured using a booking factor of 101 to 200. Under subscription (a booking value of 101 to 200) is the same as underbooking.

The term %util used in AutoRoute (AR). The %util corresponds directly to the booking factor. A %util of 10 causes 10% of configured BW for a connection to be used in the CAC algorithm. Therefore, a booking factor of 10 equals %util of 10.

A different connection BW is used for the GCAC and ACAC algorithms to deal with booking factor. For GCAC, the booking factor is considered when link BW is advertised to the network. For ACAC, the booking factor is considered when a connection is added to a link. With that in mind, here are some examples of how oversubscription and undersubscription work.

Overbooking Examples

The following examples describe different instances of overbooking/oversubscription.

Example 3-1 Booking factor is 100 on a 100 MB link

- PNNI advertises 100 MB to the network.
- The link is configured on the service module for 100MB.
- For a call of 10 MB, GCAC CACs 10 MB, and ACAC CACs 10 MB.

Example 3-2 Booking factor is 10 on a 100 MB link

- PNNI advertises 1000 MB (calculated by $100 * 100/10 = 1000$)
- The link is configured on the service module for 100MB.
- For a call of 10 MB, GCAC CACs 10MB, and ACAC CACs 1MB (using $10 * 10/100$).

Example 3-3 Booking factor is 200 on a 100 MB link

- PNNI advertises 50 MB (calculated as $100 * 100/200 = 50$)
- The link is configured on the service module for 100MB.
- For a call of 10 MB, GCAC CACs 10MB, and ACAC CACs 20 MB (using $10 * 200/100$).

PNNI advertises the effective link BW by considering the actual link BW and the booking factor. Therefore, GCAC will always CAC the configured connection BW without considering the booking factor.

The ATM interface is configured for the actual link BW, so in its ACAC, it adjusts the connection BW as required by the booking factor.

Booking factors less than 100% result in link oversubscription, since the BW booked for each connection is less than the configured BW for the connection. This is also referred to as overbooking.

Booking factors greater than 100% result in link undersubscription, and the BW booked for a connection is greater than the connection configured BW. This is also referred to as underbooking.


Note

The policing BW is still based on the configured bandwidth, not book factor. That is, for a 10MB connection, no matter what the booking factor is, the policing is still 10MB.

Infinite oversubscription is not currently supported in PNNI since the booking factor can not be set to zero (range 1 to 200).

Interface Table

The Interface Table contains topological and partition information about ports in the network. The Interface Table fields are described in Table 4-11.

Table 4-12 Interface Table (Port Specific Parameters)

Field	Defaults	Notes	Description and Comments	GET	TPI
Site	—	M/H	Site name.	*	*
PortID	0	O	Slot/port address used for linking the NMT's Bursty traffic table to the port table. Also used for bundling.	*	
FeederPort_ID	0	E	Slot/port address (cross reference) in the port table. Also used for bundling.	*	
Speed	0	O/H	Clock speed of the access port. Values range from <ul style="list-style-type: none"> • 56 to 2048 kbps for frame relay • 3622 to 38336 for ATM on an AUSM on an MGX 8220 edge concentrator. • 80000, 96000, or 353208 for ATM on a BPX, depending on the type of port Remember: FRP = Frame Relay PAD (IPX) FRSM = Frame Relay Service Module (MGX 8220/BPX) 0 = Port speed should not be changed by this record.	*	*
MinCR	0	O	Minimum Cell Rate in egress (transmit) direction for the PNNI bandwidth partition. Zero value means no partitioning.		
MaxCR	0	O	Maximum Cell rate in egress (transmit) direction for the PNNI partition. Zero value means no partitioning.		
MinLCN	0	O	Minimum number of channels in the PNNI partition. Zero value means no partitioning.		
MaxLCN	0	O	Maximum number of channels in the PNNI partition. Zero value means no partitioning.		
BF	0	O	Booking Factor used to calculate committed cell rate that contributes to the interface load. Ranges are from 1% to 200%. If 0 is specified, the globally assigned value is used for this connection. This applies to PNNI connections only, and is similar to %Util for Autoroute connections.		

